

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problems Mailbox.**



(19) Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) Publication number : 0 629 947 A2

(12)

## EUROPEAN PATENT APPLICATION

(21) Application number : 94304097.2

(51) Int. Cl.<sup>5</sup> : G06F 9/46

(22) Date of filing : 07.06.94

(30) Priority : 14.06.93 US 77211

(43) Date of publication of application :  
21.12.94 Bulletin 94/51

(84) Designated Contracting States :  
DE FR GB

(71) Applicant : International Business Machines  
Corporation  
Old Orchard Road  
Armonk, N.Y. 10504 (US)

(72) Inventor : Forman, Ira Richard  
2100 Cypress Point East  
Austin, Texas 78746 (US)  
Inventor : Madduri, Hari Haranath  
7004 Anaqua  
Austin, Texas 78750 (US)

(74) Representative : Litherland, David Peter  
IBM (UK) Ltd  
UK Intellectual Property Department  
Hursley Park  
Winchester, Hampshire SO21 2JN (GB)

### (54) Shared resource control in a distributed processing system.

(57) A system and method for determining a master process for control of a shared system resource. The improved system requires the master process to hold exclusive access on a shared resource control file only intermittently. The master process periodically updates the shared resource control file with a new timestamp. Processes seeking resource access read the shared control file and determine whether another process has been designated master. If the interval since the latest timestamp is greater than a preset staleness interval, the shared control file is discarded and a new one created by the accessing process.

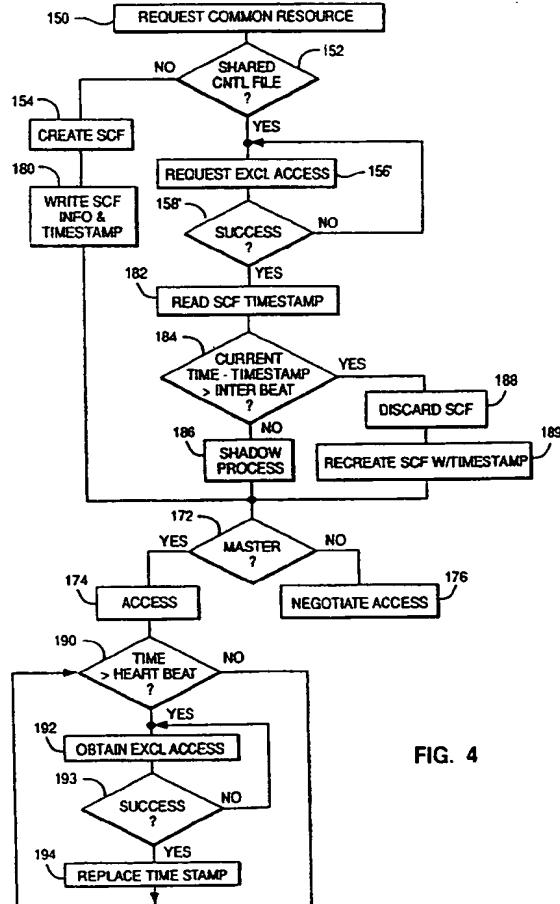


FIG. 4

The present invention relates to the operation of distributed processing computer systems. In particular it relates to those systems that have a plurality of processing nodes each one having access to a number of shared resources and which require apparatus and methods for managing access to the shared resources. Still more particularly, the present invention relates to the management of a shared control file that designates one of a number of distributed processes as the master process for controlling access to that shared resource.

Distributed computer systems are created by linking a number of computer systems using a communications network. Distributed systems frequently have the ability to share data resident on individual systems. Replicated data systems implement data sharing by providing a replica copy of a data object to each process using that data object. Replication reduces the access time for each processor by eliminating the need to send messages over the network to retrieve and supply the necessary data. A replicated object is a logical unit of data existing in one of the computer systems but physically replicated to multiple distributed computer systems. Replicated copies are typically maintained in the memories of the distributed systems.

Replicated data objects also speed the update process by allowing immediate local update of a data object. Replication introduces a control problem, however, because many copies of the data object exist. The distributed system must have some means for controlling data update to ensure that all copies of the data remain consistent.

Prior art systems control data consistency by establishing a master data object copy in one of the distributed systems. The master copy is always assumed to be valid. Data object update by a system other than that of the master copy requires sending the update request to the master for update and propagation to all replicas. This approach has the disadvantage of slowing local response time as the master data object update and propagation are performed.

Another means for controlling replicated data is described in Moving Write Lock for Replicated Objects, commonly assigned, filed on October 16, 1992 as US patent application Serial Number 07/961,757 and having attorney docket number AT992-046. This document which does not form part of the state of the art as defined in article 54(2) and 54(3) EPC describes an apparatus and method that require that a single "write lock" exist in a distributed system and be passed to each process on request. Data object updates can only be performed by the holder of the "write lock." The "write lock" holder may update the local object copy and then send that update to the master processor for its update and propagation to other processes. A copy of the above patent application is filed herewith.

The method for determining which of a number of distributed processes is to be master is described in pending US patent application serial number 07/961,750 filed October 16, 1992 and entitled Determining a Winner of a Race in a Data Processing System, commonly assigned and bearing attorney docket number AT992-117. This document does not form part of the state of the art as defined by articles 54(2) and 54(3) EPC. The "race" between each process potentially controlling a resource results in the assignment of master status to the process first establishing write control over a Share Control File. Once control has been established by one process, other processes are assigned "shadow" status. Master process failure causes reevaluation of master status. A copy of this patent application is filed herewith.

The technical problem addressed by the present invention is providing fault-tolerant features to a distributed processing system using write lock management of replicated data objects. Fault tolerance is required to ensure that no data or updates are lost due to the failure of a master process. Prior art systems and those systems referenced above, require the master determination and write lock control to be reinitialized. This could result in loss of data if a locally updated data object replica has not been propagated to the master or other replicas.

Accordingly, in a first aspect of the invention, there is provided a method of determining a master process for control of a shared resource in a computer system having a plurality of processes operating on at least one processor that has a memory and access to a shared data storage means, the method comprising: testing said shared data storage means for the presence of a shared resource control file for said shared resource; if no file exists, creating a shared resource control file in said shared data storage means, writing master process identification information to said shared resource control file, and writing a timestamp to said control file; if a file exists, requesting exclusive access to said file; if access denied, waiting and retrying; if access granted, determining the difference between current time and the last time stamp; if said difference is less than a first preset interval, designating the requesting process as a shadow process; if said difference is greater than said first preset interval, discarding said shared resource control file, creating a new shared resource control file and writing master process identification information to said shared resource control file, and writing a timestamp to said control file; if said requesting process is a master process for said shared resource, replacing the timestamp in said shared resource control file with a current timestamp after a preset second interval has passed.

Thus the present invention is directed to an improved system and method for managing a distributed processing system. The system for designating a

master process for a data object is improved to reduce the length of time exclusive control of a master process indicator is required while maintaining the ability of the indicator to show master status. The master process that wins the race for control holds temporary write access to a control indicator to update its master data and then relinquishes control. Periodically thereafter, the master process updates the control indicator. Any other process accessing the control indicator will know that another process is master unless the last master update is stale. In that case, the surviving processes again race for control.

The present invention includes a process of determining a master process for control of a shared resource in a computer system having a number of processes operating on at least one processor that has a memory and access to a shared data storage means. The process includes testing the shared data storage means for the presence of a shared resource control file for the shared resource; if no file exists, creating a shared resource control file in the shared data storage means, writing master process identification information to the shared resource control file, and writing a timestamp to the control file; if a file exists, requesting exclusive access to the file; if access denied, waiting and retrying; if access granted, determining the difference between current time and the last time stamp; if the difference is less than a first preset interval, designating the requesting process as a shadow process; if the difference is greater than the first preset interval, discarding the shared resource control file, creating a new shared resource control file and writing master process identification information to the shared resource control file, and writing a timestamp to the control file; if the requesting process is a master process for the shared resource, replacing the timestamp in the shared resource control file with a current timestamp after a preset second interval has passed.

It is an object of the present invention to improve master process efficiency by reducing the length of time exclusive control over a master process indicator is required.

A preferred embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a block diagram of a computer system of the type in which the present invention is embodied.

Figure 2 is a block diagram of a distributed network according to the present invention.

Figure 3 is a flowchart depicting the master resolution logic of an alternative system.

Figure 4 is a flowchart depicting the master resolution logic of the preferred embodiment of the present invention.

The present invention is practised in a distributed processing computer environment. This environment

consists of a number of computer processors linked together by a communications network. Alternatively, the present invention could be practised in a multi-programming system in which a single computer (e.g. single CPU) supports the execution of multiple processes each having a separate address space.

The preferred embodiment is practised with linked computers. Each computer typically has the components shown generally for the system 100 in Figure 1. Processing is provided by central processing unit or CPU 102. CPU 102 acts on instruction and data stored in random access memory 104. Long term storage is provided on one or more disks 122 operated by disk controller 120. A variety of other storage media could be employed including tape, CD-ROM, or WORM drives. Removable storage media may also be provided to store data or computer process instructions. Operators communicate with the system through I/O devices controlled by I/O controller 112. Display 114 presents data to the operator while keyboard 114 and pointing device 118 allow the operator to direct the computer system. Communications adapter 106 controls communications between this processing unit and others on a network to which it is connected by network interface 108.

Computer system 100 can be any known computer system including microcomputers, mini-computers and mainframe computers. The preferred embodiment envisions the use of computer systems such as the IBM Personal System/2 (PS/2) or IBM RISC System/6000 families of computers. (IBM, Personal System/2, PS/2 and RISC System/6000 are trademarks of the IBM Corp.) However, workstations from other vendors such as Sun Microsystems, Inc. or Hewlett Packard may be used, as well as computers from Compaq Computer Corp. or Apple Computer Corp.

A distributed processing system is shown in Figure 2. Each of the processing nodes 202, 204, 206, 208, 210 is connected to a network 200 that enables communications among the processors. Additional permanent storage may be associated with the network as shown by disk storage unit 212. In the alternative, persistent storage in one of the processing nodes could be used for network persistent storage.

Network 200 can be any type of network including LAN, WAN, ATM or other. Physical network protocols such as Ethernet or Token Ring can be used and communications protocols such as TCP/IP or Netbios or Novell Netware can control the network. Network file system management can be provided by a program based on the Sun Microsystems NFS technology or CMU AFS technology. Each of these file system programs allows distributed processes to access and manage data residing on remote systems. These systems create a single logical file system for each processor regardless of the physical location of individual files. NFS is described in greater detail in the IBM Corp. publication Communication Concepts

and Procedures, Order No. SC23-2203-00.

The variety of permitted networks means that the processing nodes may be distributed throughout a building, across a campus, or even across national boundaries.

The preferred embodiment of the present invention is practised in a distributed network of peer processing nodes. Peer nodes each have equal status in the network with none being master or slave nodes. Using peer nodes improves network efficiency because there is no single bottleneck through which requests must be funnelled. Instead each node can act independently to perform its functions. Another advantage is that failure of any particular node will not cause the entire network to fail as would be the case where a master processor existed. The disadvantage of peer networks is that there is no focal point for controlling data integrity of replicated data.

The above referenced patent application for Determining the Winner of a Race in a Data Processing System, filed herewith, teaches a procedure for "racing" for control of a resource. Figure 3 illustrates the steps of this process. The process starts by generating a request for a common resource 150. The process requesting the resource tests to determine whether or not a shared control file exists 152. If not, the process creates a shared control file 154. In either case, the process attempts to hold exclusive write access 156. If this is successful 158 the process updates the shared control file 160 and it becomes master of that resource 162. If the attempt to acquire exclusive write failed, the process is not the master 164 and must read the name of the master from the shared control file 166 and connect to the master 168 as a shadow 170. If the requesting process is the master, it can directly access the resource, otherwise, it is a shadow process and must negotiate with the master for access 176.

The shared control file of the preferred embodiment is a storage file in the logical file system. As such, it resides on one of the permanent storage devices in the distributed system. The present invention is equally applicable, however, to a shared resource control file managed in volatile memory (RAM) that is sharable among the distributed processes.

The requirement that the master process maintain an exclusive write lock 156 on the shared control file is undesirable in many systems. Each process may be master of a number of resources. Each replicated data object has a master and a particular user may cause a single process to be master of a large number of data objects. Each exclusive write lock ties up a process file descriptor. In many operating systems, the number of file descriptors allocated to each process is limited. For example, older versions of the UNIX operating system (UNIX is a registered trademark of Unix System Laboratories, Inc) allowed only four or five open file descriptors per process. Thus, a

particular process may be restricted in accessing resources because of a limit on file descriptors.

The present invention is directed to removing that limitation by allowing the master to release the exclusive write mode while still being the master process for that resource. The previous system indicated race failures to shadows by denying them the exclusive write access to the shared control file. The present invention replaces this master status indicator with a timestamp and control file age check. This change allows an unlimited number of replicated objects for each process.

The selection of the master process (the winner of the race) in the preferred embodiment of the present invention will be described with reference to Figure 4 in which reference numbers corresponding to the reference numbers of Figure 3 indicate equivalent process steps.

The process starts when a processor requests a common resource 150. The existence of a shared control file is tested 152. If no shared control file exists, the process creates one 154, obtains exclusive access, and writes identifying data including the master identity and a timestamp 180. Processing continues at test 172 where status is checked prior to access.

When a shared control file exists, the process attempts to gain exclusive write access to the file 156'. Failure to gain access means another process has exclusive access to the shared control file. In the preferred embodiment, either the master or shadow could have exclusive access to the process must retry 156' until it actually acquires exclusive access to check its status.

Success in this case does not assure the process that it is master. Instead, the process must read the timestamp value from the shared control file 182 and compare it to the current time 184. If the difference between the current time and the time stamp is less than a set period INTER\_BEAT the designated master process is still in control and the requesting process is a shadow process 186. If the difference is greater than INTER\_BEAT then the shared control file is stale. The requesting process discards the old shared control file 188 and creates a new one in which it writes its own master process identification and timestamp 189. Processing continues at step 172.

The process requesting resource access tests whether it is the master of that resource at 172. If it is the master, it may access the resource 174. If not, it is a shadow process and must negotiate for access 176. The master process must continually update the shared control file to maintain control as the master. Every HEART\_BEAT seconds 190 the master process attempts to obtain exclusive access to the shared control file. The request may fail 193 due to a shadow process holding exclusive access to check master

status. If the request fails, the master process waits and tries again. If the request succeeds, the master replaces the timestamp.

The periods INTER\_BEAT and HEART\_BEAT are set by the application developer. INTER\_BEAT - HEART\_BEAT must be greater than the expected wait time for the master process to gain exclusive write access in step 192. If it is too short, the shared control file will be prematurely invalidated. On the other hand, if INTER\_BEAT is too long transfer of control to one of the shadow processes will be delayed. The preferred embodiment sets HEART\_BEAT at 30 seconds and INTER\_BEAT at 90 seconds. These values ensure that the process does not tie up system resources in frequent timestamp updates, but provides sufficiently prompt discovery of master process failure.

The system of the present invention has the advantage of reducing the length of time a process must hold an exclusive lock. It also assures that failure of a master process can be detected and another master established without undue delay.

It will be understood from the foregoing description that various modifications and changes may be made in the preferred embodiment of the present invention. It is intended that this description is for purposes of illustration only and should not be construed in a limiting sense. The scope of this invention should be limited only by the language of the following claims.

## Claims

1. A method of determining a master process for control of a shared resource in a computer system having a plurality of processes operating on at least one processor that has a memory and access to a shared data storage means, the method comprising:

testing said shared data storage means for the presence of a shared resource control file for said shared resource;

if no file exists, creating a shared resource control file in said shared data storage means, writing master process identification information to said shared resource control file, and writing a timestamp to said control file;

if a file exists, requesting exclusive access to said file;

if access denied, waiting and retrying;

if access granted, determining the difference between current time and the last time stamp;

if said difference is less than a first preset interval, designating the requesting process as a shadow process;

if said difference is greater than said first

5 preset interval, discarding said shared resource control file, creating a new shared resource control file and writing master process identification information to said shared resource control file, and writing a timestamp to said control file;

10 if said requesting process is a master process for said shared resource, replacing the timestamp in said shared resource control file with a current timestamp after a preset second interval has passed.

15 2. A method as claimed in claim 1 wherein said shared data storage means is a network file system.

20 3. A method as claimed in claim 1 wherein said shared data storage means is shared volatile memory.

25 4. A system for controlling sharing of a plurality of shared resources in a distributed processing system having a plurality of processors connected by a communications means, the system comprising:

30 shared resource control means for storing data indicating a process as a controlling process for each of said plurality of shared resources and data indicating a time of a last update of said shared resource control means by said controlling process;

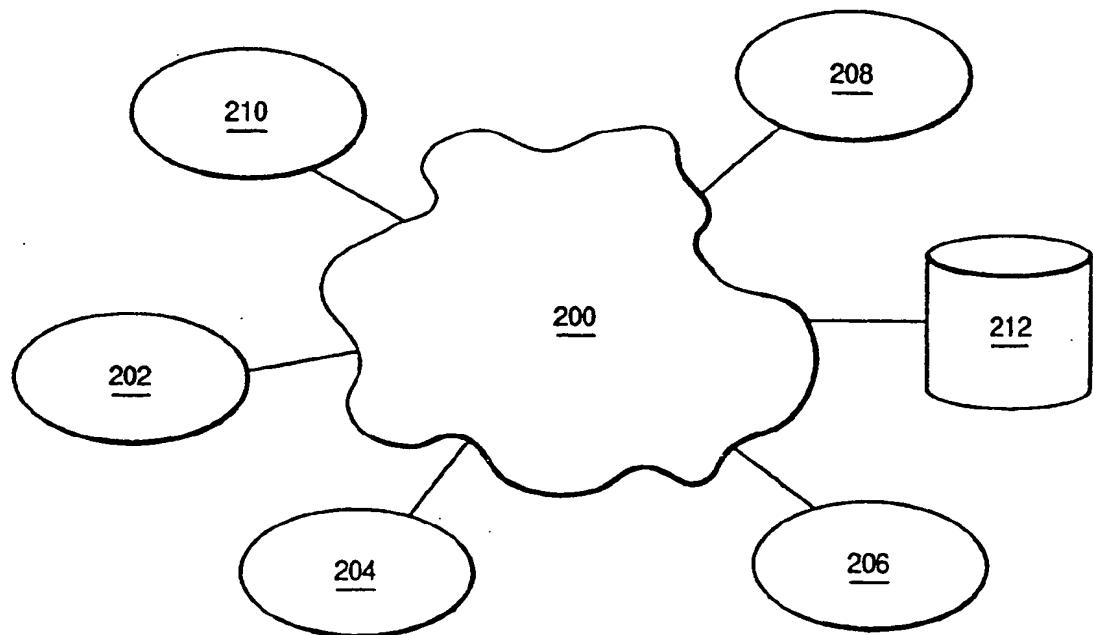
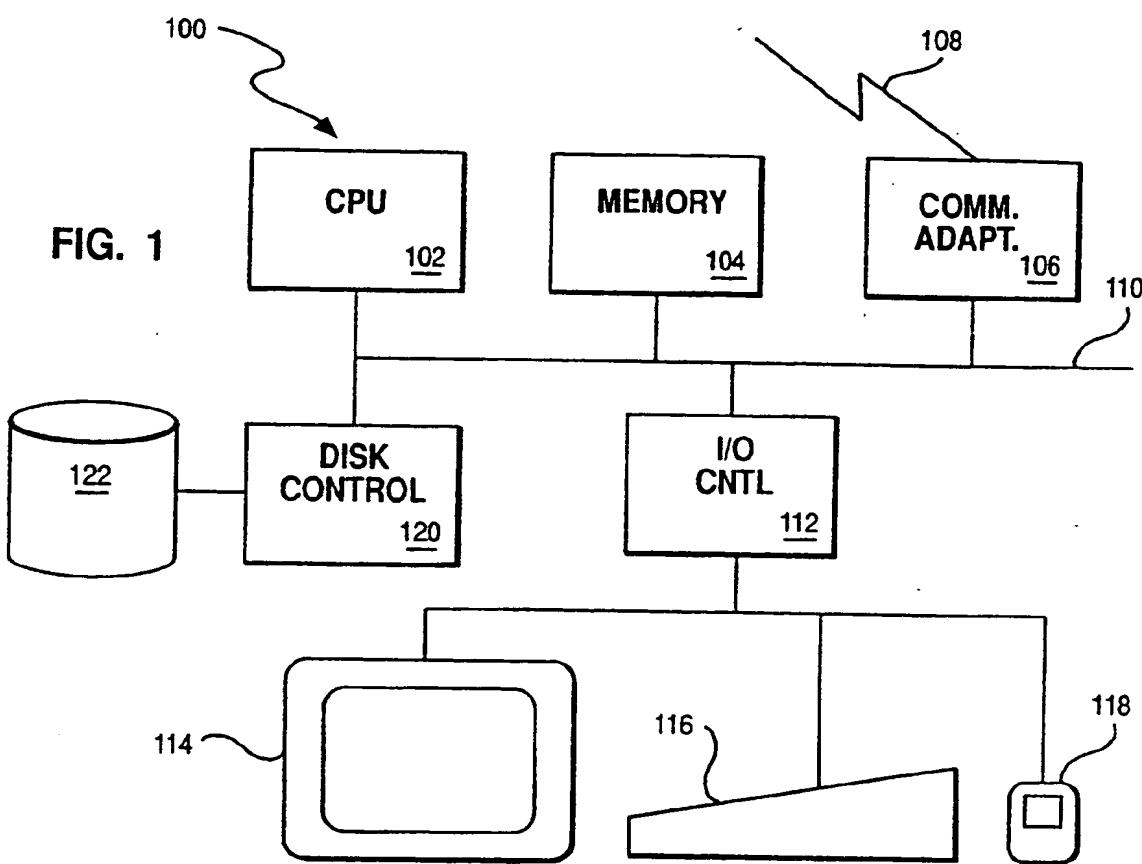
35 access control means for limiting exclusive access to said shared resource control means to at most one process;

40 shared resource control read means for reading data from said shared resource control means;

45 comparison means for comparing said data indicating the time of the last update and the current time to determine an elapsed interval;

release means for releasing said shared resource control file if said elapsed interval is greater than a first value;

50 update means for periodically causing said controlling process to replace said data indicating the time of the last update with the current time.

**FIG. 2**

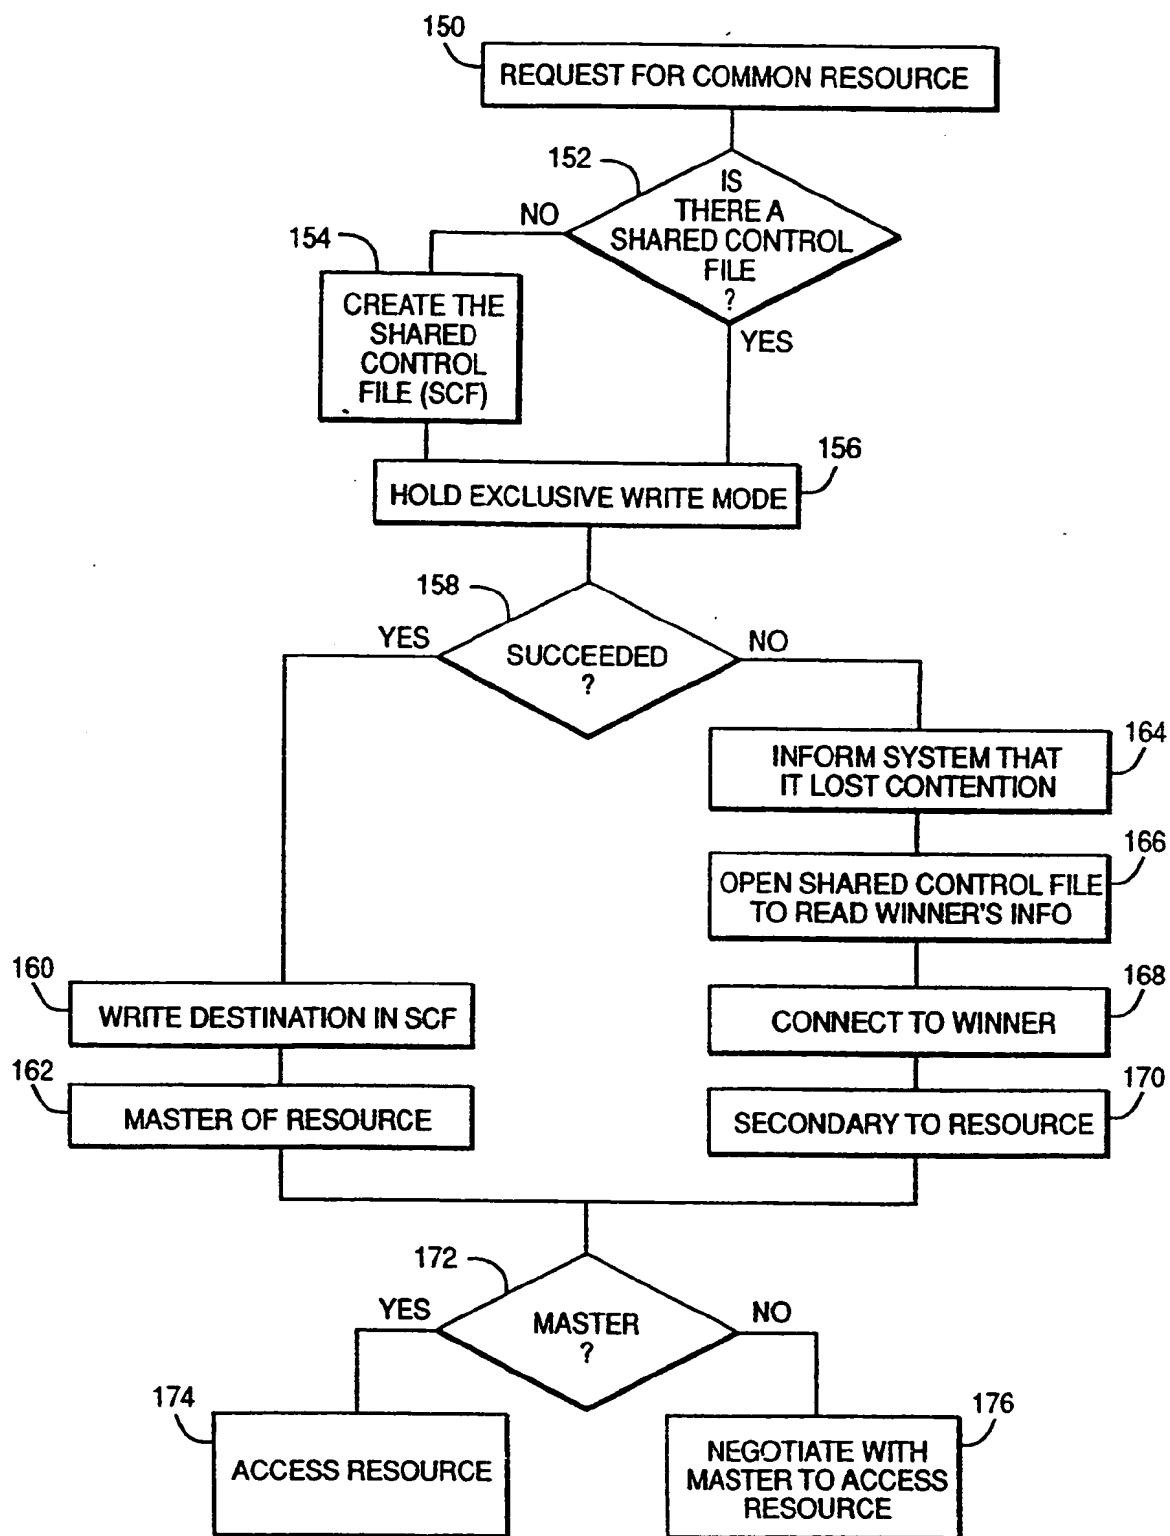


FIG. 3

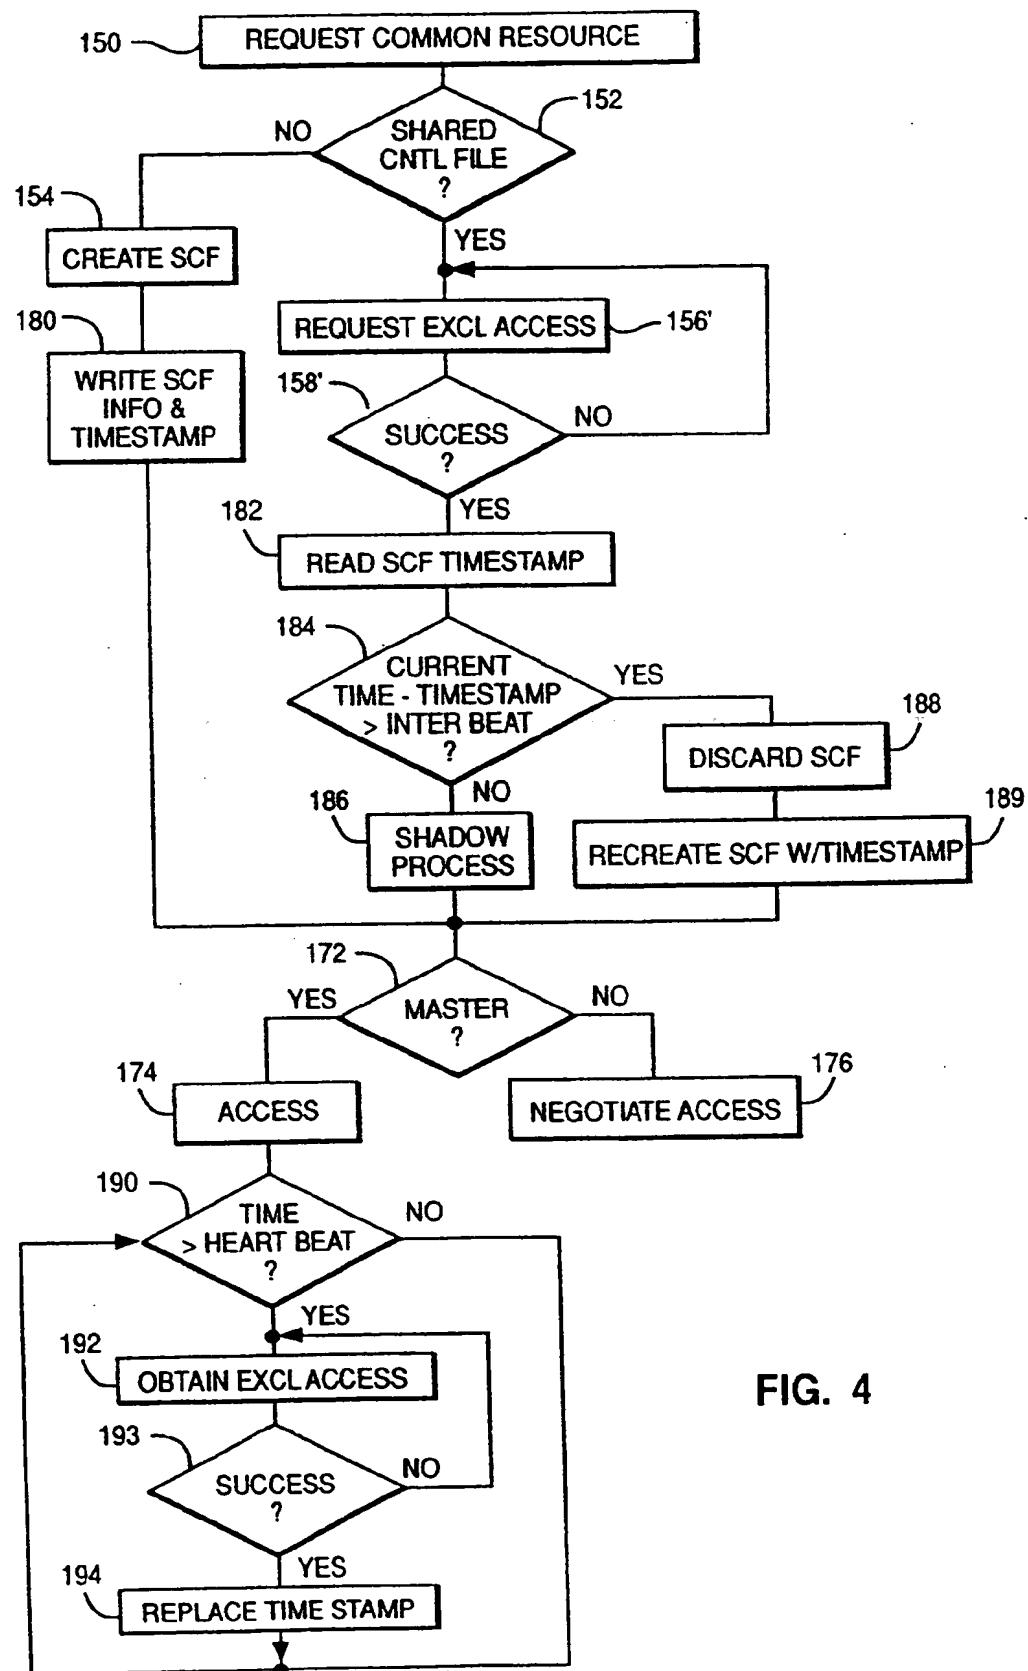
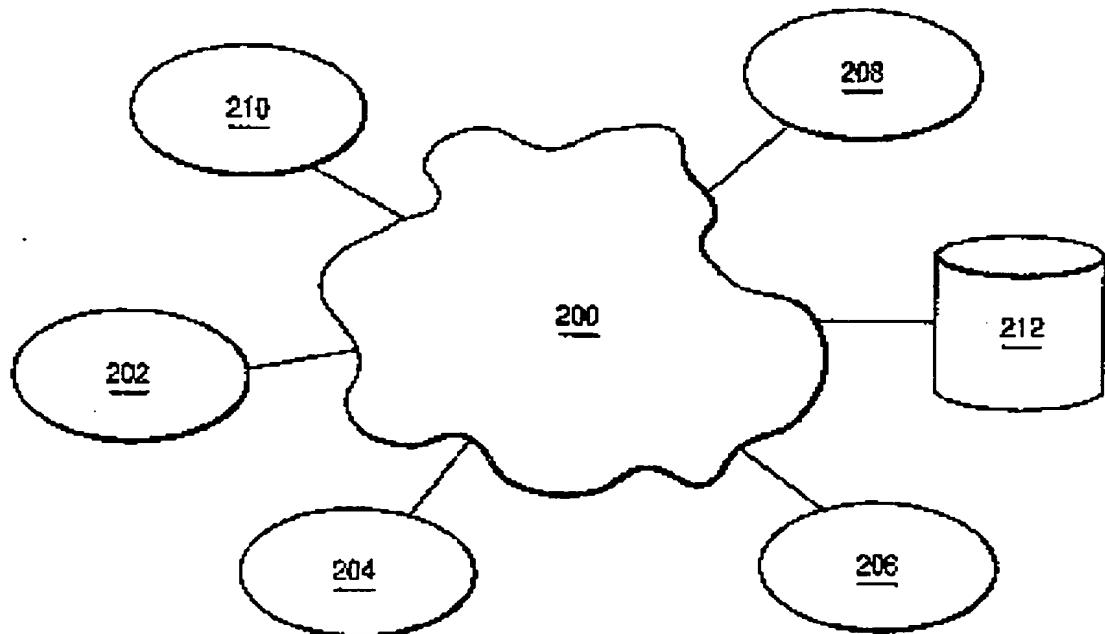
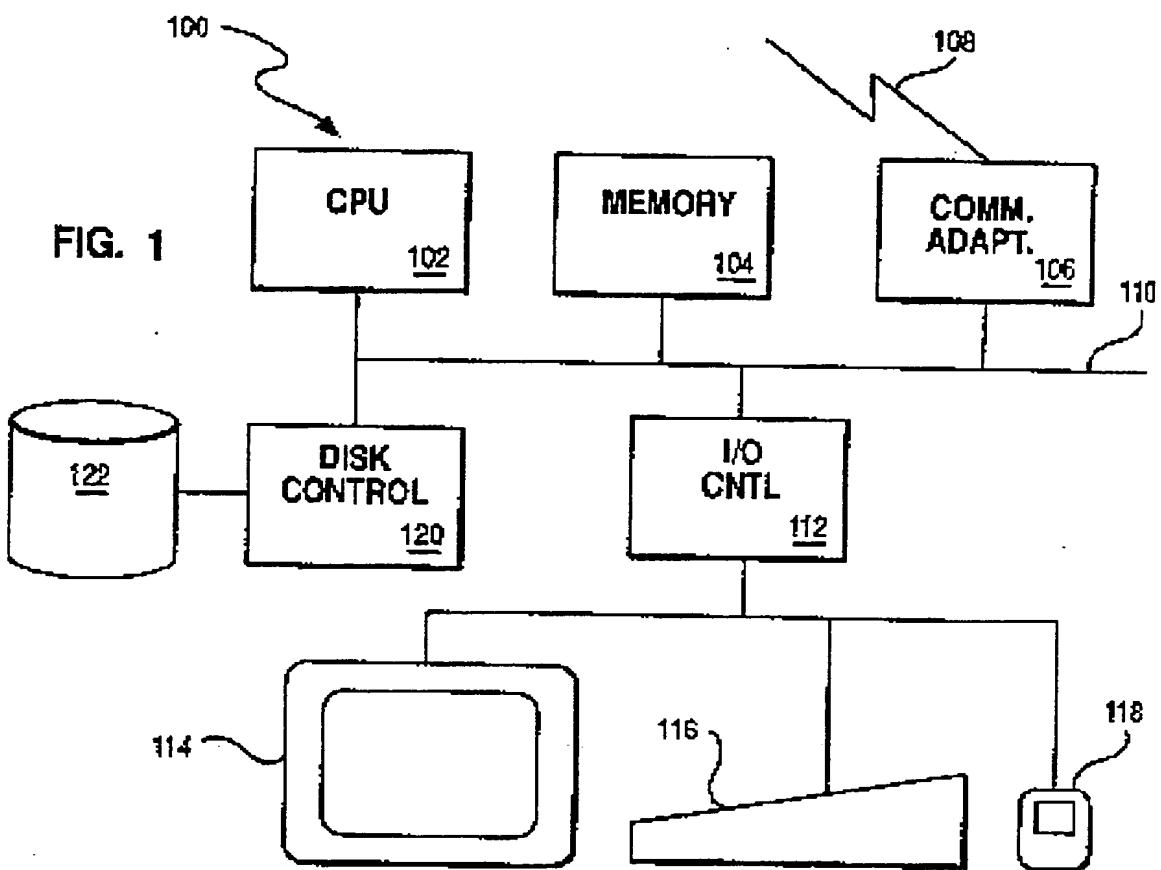


FIG. 4

**FIG. 2**

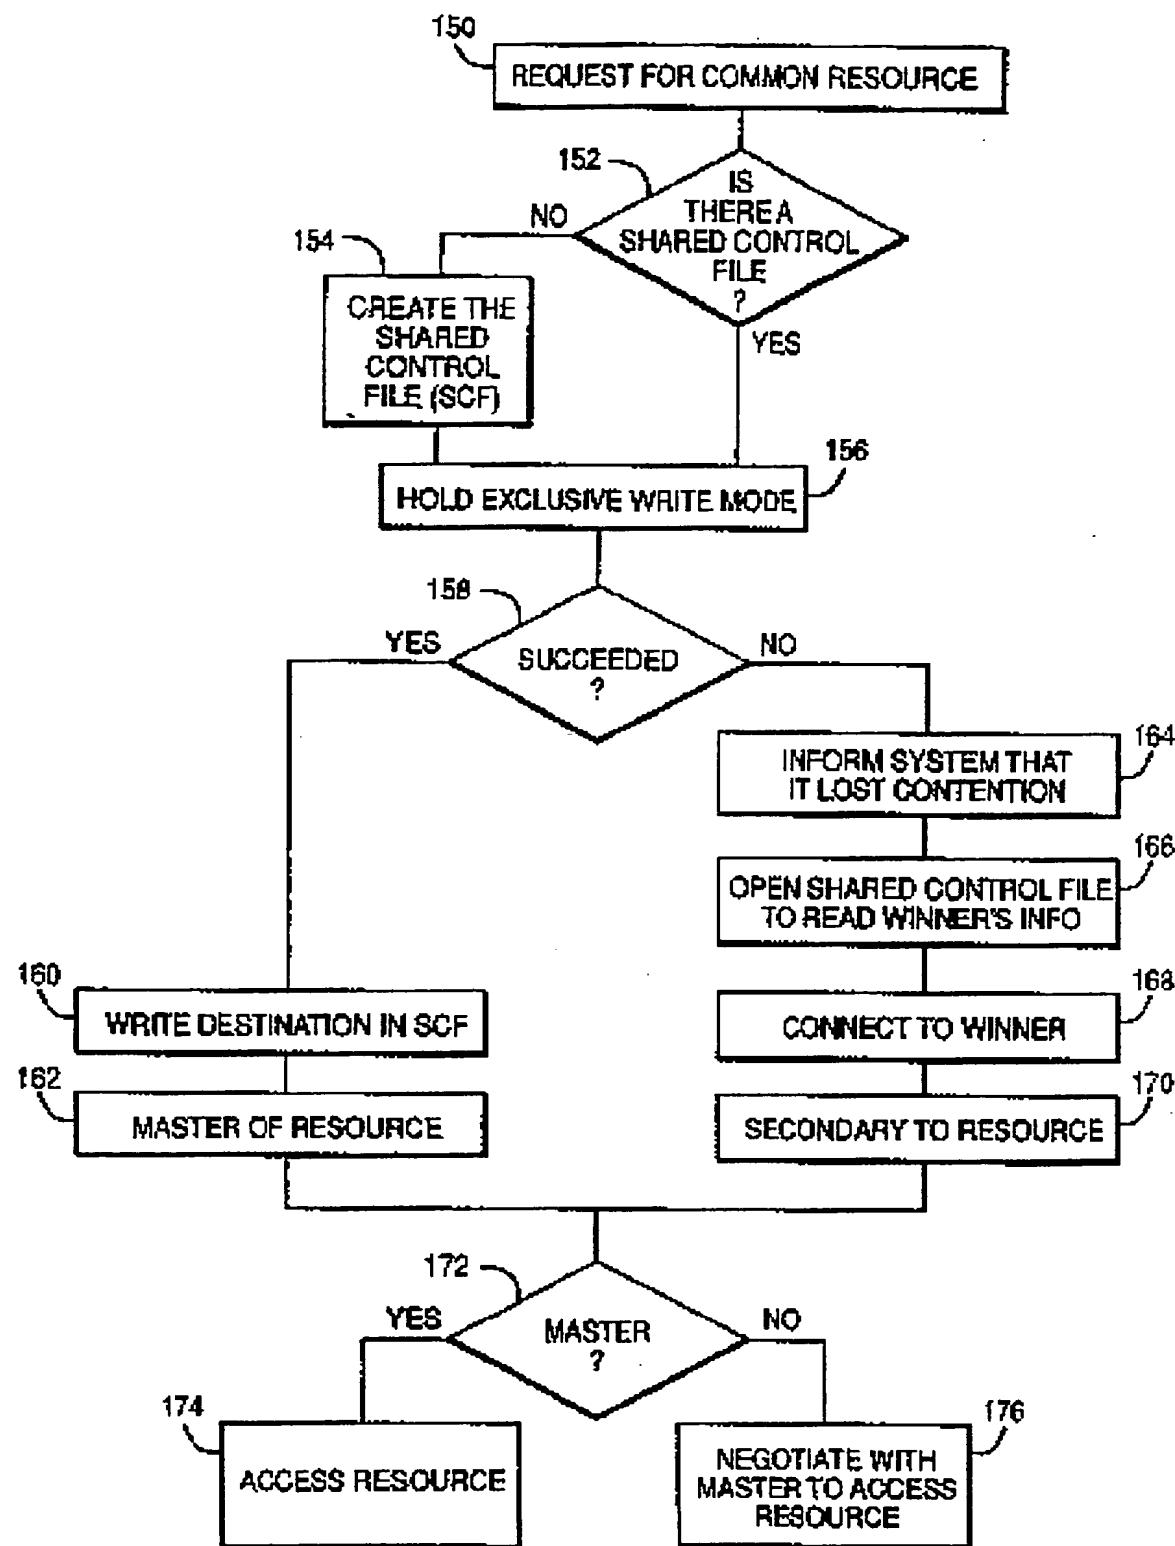


FIG. 3

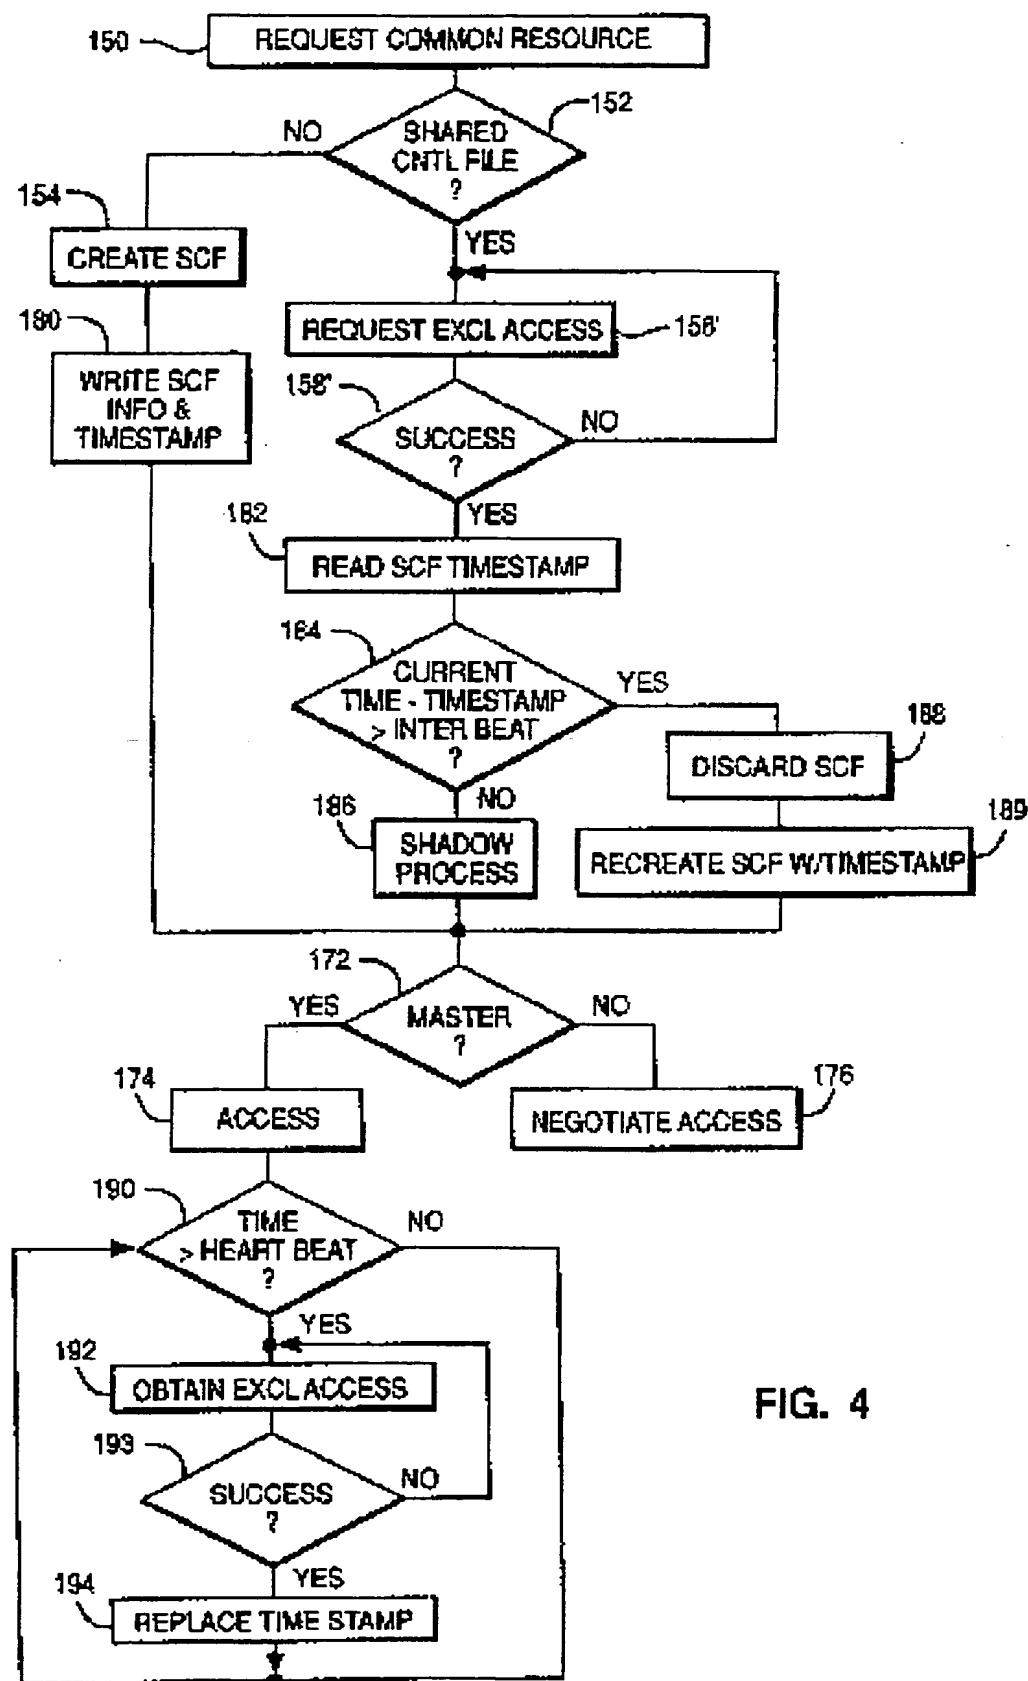


FIG. 4





Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) Publication number : 0 629 947 A3

(12)

## EUROPEAN PATENT APPLICATION

(21) Application number : 94304097.2

(51) Int. Cl.<sup>5</sup> : G06F 9/46

(22) Date of filing : 07.06.94

(30) Priority : 14.06.93 US 77211

(72) Inventor : Forman, Ira Richard

(43) Date of publication of application :  
21.12.94 Bulletin 94/51

2100 Cypress Point East  
Austin, Texas 78746 (US)  
Inventor : Madduri, Hari Haranath  
7004 Anaqua  
Austin, Texas 78750 (US)

(84) Designated Contracting States :  
DE FR GB

(74) Representative : Litherland, David Peter  
IBM United Kingdom Limited  
Intellectual Property Department  
Hursley Park  
Winchester, Hampshire SO21 2JN (GB)

(88) Date of deferred publication of search report :  
14.06.95 Bulletin 95/24

(71) Applicant : International Business Machines  
Corporation  
Old Orchard Road  
Armonk, N.Y. 10504 (US)

### (54) Shared resource control in a distributed processing system.

(57) A system and method for determining a master process for control of a shared system resource. The improved system requires the master process to hold exclusive access on a shared resource control file only intermittently. The master process periodically updates the shared resource control file with a new timestamp. Processes seeking resource access read the shared control file and determine whether another process has been designated master. If the interval since the latest timestamp is greater than a preset staleness interval, the shared control file is discarded and a new one created by the accessing process.

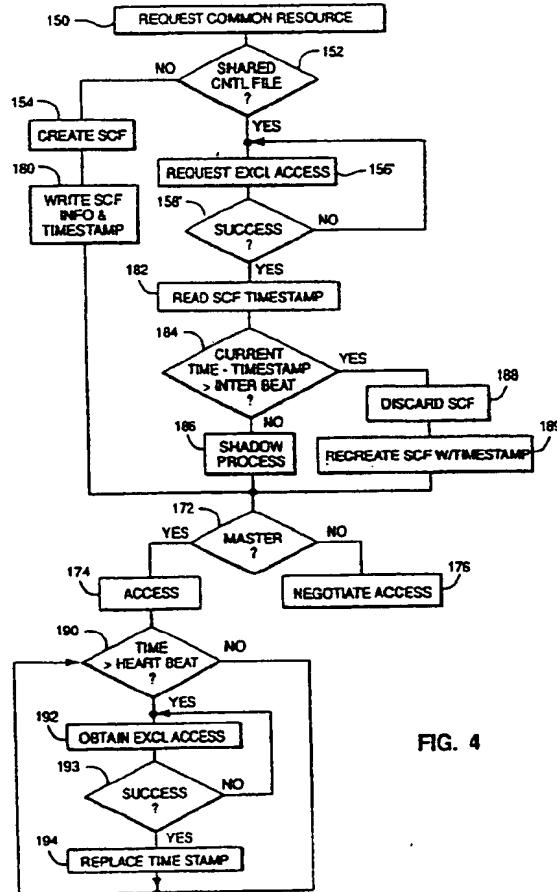


FIG. 4



European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number  
EP 94 30 4097

DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (Int.Cl.)						
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim							
A	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 34,no. 9, February 1992 NEW YORK, US, pages 248-250, ANONYMOUS 'Technique for Fast Query of a LAN Based Database.' * the whole document *	1-4	G06F9/46						
A	EP-A-0 465 871 (IBM) 15 January 1992 * the whole document *	1-4							
A	US-A-4 961 224 (YUNG DARBY) 2 October 1990 * the whole document *	1-4							
			TECHNICAL FIELDS SEARCHED (Int.Cl.)						
			G06F						
<p>The present search report has been drawn up for all claims</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%;">Place of search</td> <td style="width: 33%;">Date of completion of the search</td> <td style="width: 34%;">Examiner</td> </tr> <tr> <td>THE HAGUE</td> <td>10 April 1995</td> <td>Michel, T</td> </tr> </table>				Place of search	Date of completion of the search	Examiner	THE HAGUE	10 April 1995	Michel, T
Place of search	Date of completion of the search	Examiner							
THE HAGUE	10 April 1995	Michel, T							
<b>CATEGORY OF CITED DOCUMENTS</b> X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		I : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document							